

Control in an Information Rich World

Report of the Panel on Future Directions in
Control, Dynamics, and Systems

DRAFT: v2.1, Thu Aug 2 05:50:09 2001

Abstract

The field of *Control* provides the principles and methods used to design physical and information systems that maintain desirable performance by automatically adapting to changes in the environment. Over the last forty years, the field has seen huge advances, leveraging technology improvements in sensing and computation with breakthroughs in the underlying principles and mathematics. Automatic feedback control systems now play critical roles in many fields, including manufacturing, electronics, communications, transportation, computers and networks, and many military systems.

As we begin the 21st Century, the opportunities for, and use of, Control principles and methods is exploding. Increasingly, computation, communication and sensing will be cheap and ubiquitous, with more and more devices including embedded, high-performance processors, sensors, and networking hardware. This will make possible the development of machines with a degree of intelligence and reactivity that will change everyone's life, both in terms of the goods available and the environment in which we live.

New developments in this increasingly information rich world will require a significant expansion of the basic tool sets of Control. The complexity of the Control ideas involved in the operation of the Internet, autonomous systems, or an enterprise-wide supply chain system are on the boundary of what can be done with available methods, so new developments must be vigorously pursued.

The purpose of this report is to spell out some of the exciting prospects for the field in the current and future technological environment and to explain the critical role we expect it to play over the next decade.

Panel Membership

Richard M. Murray (chair)
California Institute of Technology

Karl Aström
Lund Inst. of Tech.

Pramod Khargonekar
University of Michigan

Stephen Boyd
Stanford University

P. R. Kumar
University of Illinois

Siva Banda
Air Force Research Laboratory

P. S. Krishnaprasad
University of Maryland

Roger Brockett
Harvard

Greg McRae
Massachusetts Inst. of Tech.

John Burns
Virginia Tech

Jerrold Marsden
California Inst. of Tech.

Munzer Dahleh
Massachusetts Inst. of Tech.

George Meyer
NASA

John Doyle
California Inst. of Tech.

William Powers
Ford

John Guckenheimer
Cornell University

Gunter Stein
Honeywell

Charles Holland
Department of Defense

Pravin Varaiya
University of California Berkeley

Additional Contributors

Raffaello D'Andrea, Hideo Mabuchi, Kevin Wise

Contents

1	Executive Summary	1
2	Overview of the Field	6
2.1	What is Control?	6
2.2	Control System Examples	11
2.3	The Shift to Information-Based Systems	16
2.4	Opportunities and Challenges Now Facing Us	18
3	Applications, Opportunities, and Challenges	24
3.1	Aerospace and Transportation	25
3.2	Information and Networks	30
3.3	Robotics and Intelligent Machines	36
3.4	Biology and Medicine	44
3.5	Materials and Processing	44
3.6	Other Applications	44
4	Education and Outreach	45
5	Recommendations	46
5.1	Integrated Control, Computation, Communications, and Net- working	46
5.2	Control at Higher Levels of Abstraction	46
5.3	High-Risk, Long-Range Applications of Control	46
5.4	Support for Theory and Interaction with Mathematics	46
5.5	New Approaches to Education and Outreach	46
5.6	Additional Considerations	46



Figure 1: Applications of Control: (a) the Watt governor, (b) flight control and (c) disk drives.

1 Executive Summary

Rapid advances in computing, communications, and sensing technology offer unprecedented opportunities for the field of Control to expand its contributions to the economic and defense needs of the nation. This report presents the conclusions and recommendations of a panel of experts chartered to examine these opportunities. We present an overview of the field, describe its successes and impact, and describe the new challenges ahead. We do not attempt to cover the entire field. Rather, we focus on those areas that are undergoing the most rapid change and that require new approaches to meet the challenges and opportunities that face us.

What is Feedback Control?

At its simplest, a control system represents a feedback loop in which a sensed quantity is used to modify the behavior of a system through computation and actuation. Control systems engineering traces its roots to the industrial revolution, to devices such as the Watt flyball governor, shown in Figure 1. This device used a flyball mechanism to sense the rotational speed of a steam turbine, adjusting the flow of steam into the machine using series of linkages. By thus regulating the turbine's speed, it provided the safe, reliable, consistent operation that was required to enable the rapid spread of steam-powered factories.

A more modern view of Control sees feedback as a tool for uncertainty management. By measuring the operation of a system, comparing it to a

reference, and adjusting available control variables, we can cause the system to respond properly even if its dynamic behavior is not exactly known or if external disturbances tend to cause it to respond incorrectly. It is precisely this aspect of Control as a means for ensuring robustness to uncertainty that explains the fact that feedback control systems are all around us in the modern technological world. They are in our homes, cars and toys, in our factories and communications systems, and in our transportation, military and space systems.

The use of Control is extremely broad, but can be roughly broken down into three basic categories:

- Control of electromechanical systems, where computer-controlled actuators and sensors regulate the behavior of the system
- Control of electronic systems, where feedback is used to compensate for component variations and provide reliable, repeatable performance
- Control of information and decision systems, where limited resources are dynamically allocated based on estimates of future needs.

The contributions to the field of Control come from many disciplines, including pure and applied mathematics; aerospace, chemical, mechanical, and electrical engineering; operations research and economics; and the physical and biological sciences.

Why Does It Matter?

Over the past 40 years, the advent of analog and digital electronics has allowed Control technology to spread far beyond its initial applications, and has made it an enabling technology in many applications. Visible successes from past investment in Control include:

- Guidance and control systems for aerospace vehicles, including commercial aircraft, guided missiles, advanced fighter aircraft, launch vehicles, and satellites. All provide stability and tracking in face of environmental and system uncertainties.
- Control systems in the manufacturing industries, from automotive to integrated circuits. In the latter, the positioning stages of IC steppers rapidly position and align wafers to the extraordinary degree of accuracy required for high yield fabrication.



Figure 2: Modern networked systems: (a) the California power network, (b) UUNET's North American backbone and (c) Chemotaxis in *E. Coli*.

- Industrial process control systems, particularly in the hydrocarbon and chemical processing industries, that maintain high product quality by monitoring hundreds of sensor signals and making corresponding adjustments to hundreds of valves, heaters, pumps, and other actuators.

These applications have had an enormous impact on the productivity of modern society.

In addition to its impact on modern engineering applications, Control has also made significant intellectual contributions. Control theorists and engineers have made rigorous use of and contributions to mathematics, motivated by the need to develop provably correct techniques for design of feedback systems. They have been consistent advocates of the “systems perspective,” and have developed reliable techniques for modeling, analysis, design, and testing that enable development and implementation of the wide variety of very complex engineering systems in use today. Moreover, the Control community has been a major source and training ground for people who embrace this systems perspective and who master the substantial set of knowledge and skills it entails.

Control Will Be Even More Important in the Future

As we look forward, the opportunities for new applications and new advances in Control expand dramatically. The advent of ubiquitous, distributed computation, communication, and sensing systems has begun to create an environment in which we have access to enormous amounts of data, and the ability to process and communicate that data in ways that were unimagined 20 years ago. This will have a profound effect on Control, especially as our software systems begin to interact with our physical systems in much more integrated ways. Figure 2 illustrates several systems where these trends are evident.

In all of these areas, a common feature is that system level requirements far exceed the achievable reliability of individual components. This is precisely where feedback (in its most general sense) plays a central role, since it allows the system to ensure that it is achieving its goal, through correction of its actions based on sensing its current state. The challenge to the field is to go from the traditional view of control systems as an interconnected set of components, to realizing control systems as a heterogeneous collection of physical and information systems, with intricate interconnections and interactions.

In addition to cheap and ubiquitous computation, communication and sensing, and the shift from physics-based systems to information-based systems, an important trend in Control is the move from low-level control to higher levels of abstraction. This includes such things as increased autonomy in flight systems (all the way to complete unmanned operation), integration of local feedback loops into enterprise-wide scheduling and resource allocation systems, and control of systems with linguistic and symbolic descriptions. Extending the benefits of Control to these larger scale systems offers enormous opportunities in improved efficiency, productivity, safety, and reliability.

... But It Won't Be Easy

In order to realize the potential of Control applied to these new applications, new methods and approaches must be developed. Among the challenges currently facing the field, a few examples provide insight into the difficulties ahead:

- Control in distributed, asynchronous, networked environments
- High level coordination and autonomy
- Automatic synthesis of control laws, with integrated validation and verification
- Building very reliable systems from unreliable parts.

Each of these challenges will require many years of effort by some of the best minds in the field to make the results rigorous, practical, and widely available. They will require investments by funding agencies to insure that our current progress is continued and that the opportunities are realized.

What Needs to Be Done

We recommend that the following actions be undertaken to address these challenges and deliver on the promise of the Control field:

1. Substantially increase research aimed at the *integration* of control, computer science, communications, and networking. This includes principles, methods and tools for control of high level, networked, distributed systems, and rigorous techniques for reliable, embedded, real-time software.
2. Substantially increase research in Control at higher levels of abstraction, moving toward enterprise level systems. This includes work in dynamic resource allocation in the presence of uncertainty, and learning, adaptation, and artificial intelligence for dynamic systems.
3. Explore high-risk, long-range applications of Control to areas such as nanotechnology, quantum mechanics, biology, and environmental science. Dual investigator funding might be particularly useful mechanism in this context.
4. Maintain support for theory and interaction with mathematics, broadly interpreted (including areas such as dynamical systems, graph theory, combinatorics, complexity theory, queuing theory, etc). The strength of the field relies on its close contact with rigorous mathematics, and this will be increasingly important in the future.
5. Invest in new approaches to education and outreach for the dissemination of basic ideas to non-traditional audiences. For Control to realize its potential, we must do a better job of educating a broader range of scientists and engineers on the principles of feedback and its use as a tool for altering the dynamics of systems and managing uncertainty.

These actions build upon the rich heritage of rigorous work in Control, extending that work to cover ever more complex and significant technological problems. They are key actions to realize the opportunities of Control in the future information-rich world.

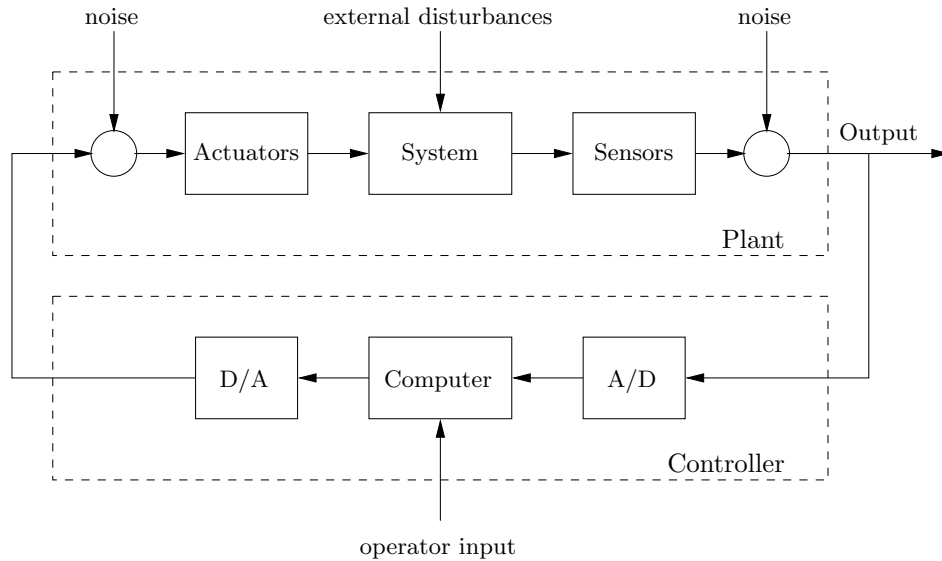


Figure 3: Components of a modern control system.

2 Overview of the Field

Control is a field with broad application to a number of engineering applications. Its impact on modern society is both profound and poorly understood. In this chapter, we provide an overview of the field, illustrated with examples and vignettes, and describe the new environment for Control.

2.1 What is Control?

The term “control” has many meanings and often varies between communities. In this report, we define Control to be the use of algorithms and feedback in *engineered* systems. Thus, Control includes such examples as feedback loops in electronic amplifiers, proportional integral derivative (PID) controllers in chemical processing, “fly by wire” systems on aircraft, and even router protocols that control traffic flow on the Internet. At its core, Control is an *information* science, and includes the use of information in both analog and digital representations.

A typical example of a modern control system is shown in Figure 3. The basic elements of a automatic control system are *sensing*, *actuation* and *computation*. In modern feedback systems, computation is typically implemented on a digital computer, requiring the use of analog-to-digital (A/D) and digital-to-analog (D/A) converters. *Uncertainty* enters the sys-




Figure 4: Feedback in Biological Systems.

tem through noise in sensing and actuation subsystems, external disturbances that affect the underlying system physics, and uncertain dynamics (parameter errors, unmodelled effects, etc). Through careful design of the feedback system, the overall dynamics of the plant can be modified to give a desirable response to operator input and the effects of uncertainty can be minimized.

It is important to note that while feedback is a central element of controls engineering, feedback as a phenomenon is ubiquitous in science and nature. Biological systems are one example. Figure 4 shows a schematic diagram of the sensing, actuation, and computation responsible for chemotaxis in a bacterium. In this system, receptors (sensors) detect gradients in stimulant concentration, chemical and electrical networks modulate “run” versus “tumble” commands (computation) and the flagella generates the desired motion (actuation).

While ideas and tools from control theory can be applied to these systems, we focus our attention in this report on the application of feedback to *engineering systems*. We also limit ourselves to a small subset of the many aspects of Control, choosing to focus on those that are undergoing the most change are most in need of new ideas and new techniques.

Control Theory

Over the last 50 years, there has been careful attention by control theorists to the issues of completeness and correctness. This includes substantial efforts by mathematicians and engineers to develop a solid foundation for proving stability and robustness of feedback controlled systems, and the development of computational tools that provide guarantees in performance in the presence of uncertainty. This rigor in approach is a hallmark of modern control theory and is largely responsible for the success it has enjoyed across

a variety of disciplines.

It is useful in this context to provide a brief history of the development of modern control theory.

Automatic control traces its roots to the beginning of the industrial revolution, when simple governors were used to automatically maintain steam engine speed despite changes in loads, steam supply, and equipment. In the early 20th Century, the same principles were applied in the emerging field of electronics, yielding feedback amplifiers that automatically maintained constant performance despite large variations in vacuum tube devices.

The development of control theory began in the 1940s, with the development of methods for single-input, single-output feedback loops, including transfer functions and Bode plots for modeling and analyzing frequency response and stability, and Nyquist plots and gain/phase margin for studying stability of feedback systems. This first generation of techniques is known collectively as “classical control” and is still the standard introduction to controls for engineering students.

In the 1960s, the second generation of control theory, known as “modern control,” was developed to provide methods for multi-variable systems where many strongly coupled loops must be designed simultaneously. These tools made use of state space representations of control systems and were coupled with advances in numerical optimization and optimal control.

Finally, the third generation of control theory, known as “robust multi-variable control”, added powerful formal methods to guarantee desired closed loop properties in the face of uncertainties. In many ways, robust control brought back some of the key ideas from the early theory of control, where uncertainty was a dominant factor in the design methodology. Techniques from operator theory were extremely useful here and there was stronger interaction with mathematics, both in terms of using existing techniques and developing new mathematics.

Over the past two decades, many other branches of control have appeared, including adaptive, nonlinear, geometric, hybrid, fuzzy, and neural control frameworks. All of these have built on the tradition of linking applications, theory and computation to develop practical techniques with rigorous mathematics.

A closely related area is the problem of dynamic allocation of resources under uncertainty, which arises in many problems in operations research. Although not normally considered part of Control, at least in the traditional sense, it relies on many of the same underlying tools (e.g., optimization and game theory) and is an increasingly important aspect as control applications move from physics-based systems to information-based systems.

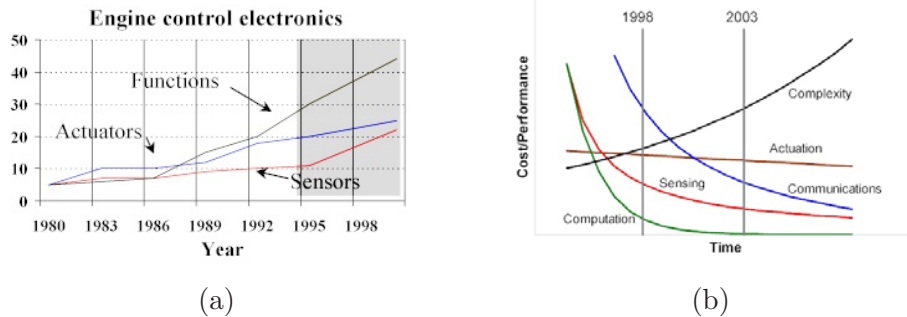


Figure 5: Trends in Control Technology.

Control Technology

Control *technology* includes modeling, sensing, actuation and computation, used together to produce a working system. A modern control loop senses the operation of a system, compares that against the desired behavior, computes corrective actions based on a model of the system's response to external inputs, and actuates the system to effect the desired change. The key issues in designing control logic are insuring that the dynamics of the closed loop plant are stable (bounded disturbances give bounded errors) and that dynamics have the desired behavior (good disturbance rejection, fast responsiveness to changes in operating point, etc). These properties are established using a variety of modeling and analysis techniques that capture the essential physics of the system and permit the exploration of possible behaviors in the presence of uncertainty, noise, and component failures.

Figure 5(a) shows some of the trends in sensing, actuation, computation and communications in automotive applications. As in many other application areas, the number of sensors, actuators, and microprocessors is increasing dramatically, as new features such as anti-lock brakes, adaptive cruise control, active restraint systems, and enhanced engine controls are brought to market. The cost/performance curves for these technologies, as illustrated in Figure 5(b), is also insightful. The costs of electronics technologies, such as sensing, computation, and communications, is decreasing dramatically, enabling more information processing. Perhaps most important is the role of communications, which is now cheap enough and ubiquitous enough to offer many new possibilities.

Control is also closely related to the integration of software into physical systems. Virtually all modern control systems are implemented using digital computers. Often they are just a small part of much larger computing

systems performing various other system management tasks. Because of this, control software becomes an integral part of the system design and is an enabler for many new features in products and processes. Online reconfiguration is a fundamental feature of computer controlled systems and this is, at its heart, a control issue.

Feedback as a Tool

Feedback is a powerful tool in making engineering systems behave in a desired fashion. Through feedback, we can *alter the dynamics* of a system to meet the needs of an application. Thus, systems that are unstable can be stabilized, systems that are sluggish can be made responsive, and systems that have drifting operating points can be held constant. In addition, feedback can be used to *manage uncertainty*. By measuring the operation of the system and comparing it to a reference, we can force the system to respond properly even if the systems dynamics are not exactly known or external disturbances are present that would normally force the system to respond incorrectly.

One of the powerful features of modern control theory is that it provides an *explicit* framework for representing uncertainty. Thus, we can describe a “set” of systems that represent the possible instantiations of a system or the possible descriptions of the system as it changes over time. While this framework is important for all of engineering, the Control community has developed one of the most powerful collection of tools for dealing with uncertainty. This was necessary because the use of feedback is not entirely benign. In fact, it can lead to catastrophic failure if the uncertainty is not properly managed (through positive feedback, for example).

Control also provides new techniques for (control-oriented) system modeling and identification. Since models play an essential role in analysis and design of feedback systems, sophisticated tools have been developed to build such models. These include input/output representations of systems (how do disturbances propagate through the system) and data-driven system identification techniques. The use of “forced response” experiments to build models of systems is well developed in the control field and these tools find application in many disciplines, independent of the use of feedback.

[Note]

Add 1–2 paragraphs about growth of embedded systems and the role this will play in active alternation of dynamic behavior and uncertainty management. Key idea is that this is now much more practical than it was. There are some good examples from the automotive industry that we might use

here.

Comparison with Physics, Computer Science, and Other Disciplines

Control engineering relies on and shares tools from physics (dynamics and modeling), computer science (information and software) and operations research (optimization and game theory), but it is also very different from these subjects, in both insights and approach.

Modeling of physical systems is common across all engineering and scientific disciplines. One of the key differences in control-oriented modeling is the way in which interactions between subsystems (components) are represented. Control relies on input/output modeling that allows many new insights into the behavior of systems by capturing the forced response of systems (to actuation and disturbances) and allowing coupling between subsystems. For example, disturbance rejection is naturally described by input/output models: we wish to design controllers that minimize the effects of disturbances (inputs) on the regulated variables of the system (outputs). Model reduction, where a simpler (lower-fidelity) description of the dynamics is derived from a high fidelity model, is also very naturally described in an input/output framework. Perhaps most importantly, modeling in a control context allows us to design *robust* interconnections between subsystems, a feature that is crucial in the operation of all large, engineered systems.

Similarly, although many control algorithms are implemented in software, these algorithms and software are very different from traditional computer software. The physics (dynamics) of the system are paramount in analyzing and designing them and their (hard) real-time nature dominates issues of their implementation. From a software-centric perspective, an F-16 is simply another peripheral, while from a control-centric perspective, the computer is just another implementation medium for the feedback law. Neither of these are adequate abstractions, and this is one of the key areas identified in this report as both an opportunity and a need.

2.2 Control System Examples

Control systems are all around us in the modern technological world. They maintain the environment, lighting, and power in our buildings and factories, they regulate the operation of our cars, toys, and manufacturing processes, they enable our transportation and communications systems, and they are critical elements in our military and space systems. For the most part,

they are hidden from view, buried within the code of processors, executing their functions accurately and reliably. Nevertheless, their existence is a major intellectual and engineering accomplishment that is still evolving and growing, promising ever more important consequences to society.

Early Examples

The proliferation of control in engineered systems has occurred primarily in the latter half of the 20th Century. There are some familiar exceptions, of course, such as the Watt governor, described earlier, and the thermostat, designed at the turn of the century to regulate temperature of buildings.

The thermostat, in particular, is often cited as a simple example of feedback control that everyone can understand. Namely, the device measures the temperature in a building, compares that temperature to a desired set point, and uses the “feedback error” between these two to operate the heating plant, e.g., to turn the plant “on” when the error is negative and “off” again as it grows positive. This explanation captures the essence of feedback, but it is a bit too simple even for a basic device such as the thermostat. Actually, because lags and delays exist in the heating plant and sensor, a good thermostat does a bit of anticipation, turning the plant off before the error actually changes sign. This avoids excessive temperature swings and cycling of the heating plant.

This modification illustrates that, even in simple cases, good control system design is not entirely trivial. It must take into account the dynamic behavior of the object being controlled in order to do a good job. The more complex the dynamic behavior is, the more elaborate the modifications must be. In fact, the development of a thorough theoretical understanding of the relationship between dynamic behavior and good controllers constitutes the most significant intellectual accomplishment of the Control community, and the codification of this understanding into powerful computer aided engineering design tools makes all modern control systems possible.

There are many other control system examples, of course, that have developed over the years with progressively increasing levels of sophistication and impact. An early system with very broad public exposure was the “cruise control” option introduced on automobiles in the mid-1960s. With cruise control, ordinary people experienced the dynamic behavior of closed loop feedback systems in action—the slowdown error as the system climbs a grade, the gradual reduction of that error due to integral action in the controller, the small (but unavoidable) overshoot at the top of the climb, etc. More importantly, by experiencing these systems operating reliably

and robustly, the public learned to trust and accept feedback systems, permitting their increasing proliferation all around us. Later control systems on automobiles have had more concrete impact, of course, such as emission controls and fuel metering systems that have achieved major reductions of pollutants and increases in fuel economy.

In the industrial world, control systems have been key enabling technologies for everything from factory automation, starting with numerically controlled machine tools, to process controls in oil refineries and chemical plants, to IC manufacturing, and to power generation and distribution over the power grid. They now also play critical roles in the routing of messages across the Internet (TCP/IP) and in power management on wireless communication systems.

Aerospace and Flight Control

Similarly, control systems have been critical enablers in the aerospace and military world. We are all familiar, for example, with the saturation bombing campaigns of World War II, needing to drop unguided explosives almost indiscriminately on population centers in order to destroy selected industrial or military targets. These have been replaced with precision guided weapons with uncanny accuracy, a single round for a single target. This is enabled by very sophisticated control systems, combining inertial guidance sensors, radar and infrared homing seekers, satellite navigation updates from the global positioning system, and very sophisticated processing of the “feedback error,” all combined in an affordably disposable package.

We are also all familiar with early space launches. Slender rockets balanced precariously on the launch pad, failing too often in out-of-control tumbles or fireballs shortly after ignition. Robust, reliable, and well-designed control systems are not optional here, because boosters themselves are unstable. And control systems have lived up to this challenge. We now take routine launch operations for granted, supporting manned space stations, probes to the outer planets, and a host of satellites for communications, navigation, surveillance, and earth observation missions. Of course, these payloads are themselves critically dependent on robust, reliable and well-designed control systems for everything from attitude control, on-orbit station-keeping, thermal management, momentum management, communications, etc.

Another notable success story for control in the aerospace world comes from the control of flight. More dramatically than many others, this example illustrates just how significant the intellectual and technological accomplish-

ments of control have been and how important their continued evolution will be in the future.

Control has played a key role in the development of aircraft from the very beginning. Indeed, the Wright brother's first powered flight was successful only because the aircraft included control surfaces (warpable wings and forward-mounted vertical and horizontal fins) that were adjusted continuously by the pilot to stabilize the flight. These adjustments were critical because the Wright Flyer itself was unstable, unable to maintain steady flight on its own.

Because pilot workload is high when flying unstable aircraft, most early aircraft that followed the Wright Flyer were designed to be statically stable. Still, as the size and performance capabilities of aircraft grew, their handling characteristics deteriorated. Designers then installed so-called "stability augmentation systems"—automatic control systems designed to modify dynamic behavior of aircraft slightly in order to make them easier to fly. These systems first appeared during the World War II years. They used early inertial sensors to measure flight motions, early analog electronic systems to construct and process feedback errors, and early hydraulic systems to actuate the linkages of selected control surfaces (vertical and horizontal tails, ailerons, etc).

Two issues surfaced immediately as these systems were being fielded: (1) how to design the control logic systematically (early systems were essentially developed by trial-and-error), and (2) how to build the systems such that they would operate reliably. (Early systems proved to be quite unreliable. Hence, only a small fraction of the full authority of the control surfaces was typically allocated to the automatic system, with the bulk of authority reserved for manual control, so the pilot could always override the automation.)

Control theorists provided the solution for the first issue. They developed modeling and simulation methods (based on differential equations and transfer functions) that accurately describe aircraft dynamics, and they developed increasingly powerful generations of control analysis and design methods to design control laws. Classical control methods enabled the systematic design of early stability augmentation systems, while modern control and robust multi-variable control are critical in all of today's modern flight systems.

But analysis and design methods alone could not address the second issue of early stability augmentation systems, namely the need for highly reliable control implementations. That issue was resolved with the development of airborne digital computers and redundant architectures. These are now

routinely used on all commercial and military aircraft. They have become so highly reliable that the old solution of granting only partial authority to automation has long been abandoned. In fact, most modern flight control implementations do not even include mechanical linkages between pilots and control surfaces. All sensed signals and control commands go through the digital implementation (e.g., fly-by-wire).

Today, we even entrust the very survival of aircraft to the automation. Examples include the all weather auto-land functions of commercial transports, in which safe go-around maneuvers are not available if failures were to occur at certain critical flight phases. Other examples include the F-16, B-2, and X-29 military aircraft, whose basic dynamics are unstable, like the Wright Flyer, but so much more violently that manual stabilization is not possible. Finally, in modern flight systems there is a growing trend to automate more and more functions—all the way to removing the pilot entirely from the cockpit. This is already commonplace in certain military reconnaissance and surveillance missions and will soon be extended to more lethal ones, such as suppressing enemy air defenses with unmanned air vehicles (UAVs).

Vignette: Fighter Aircraft and Missiles (Kevin A. Wise, Boeing Corporation)

The 1990s has been a decade of significant accomplishments and change for the aerospace community. New systems such as unstable, tailless aircraft, propulsion controlled ejection seats, and low-cost, accurate, GPS guided munitions were developed. Fly-by-wire flight control systems have become the standard, making control system design and analysis central to military aircraft and missile system development. Improving pilot safety and reducing costs were key focus areas in industry.

Flight control system design methods using feedback linearization paved the way for new gain scheduled flight control systems for aircraft. This method, applied to the X-36 Tailless Agility Research aircraft and the F-15 ACTIVE, uniquely allows engineers to better design flying qualities into the aircraft, reducing design and development costs and improving pilot acceptance. Advances in robustness theory improved analysis tools allowing engineers to accurately predict and thus expand departure boundaries for these highly unstable aircraft. To further improve safety, these control laws were augmented with neural networks for reconfigurable and damage adaptive flight control.

Missile systems, such as the Joint Direct Attack Munition (JDAM) and the Miniaturized Munition Technology Demonstrator (MMTD) developed their flight control designs using state feedback optimal control, and then projecting out those states not measured by sensors. This method eliminated sensor hardware, reducing weight and costs, and proved to be completely automatable. The Fourth Generation Escape System (GEN4) ejection seat also used this approach for its control laws. In addition to needing optimal performance, advances in robustness theory were used to characterize the seat's control system performance to uncertain crew member size and weight (95% male to 5% female). Autocode software tools for implementing controls systems also emerged in the 1990s. These computer aided design tools provide a single environment for control design and analysis as well as software design and test. They have greatly reduced the implementation and testing costs of flight control systems.

The new challenge faced by the control community is the development of unmanned combat systems (munitions as well as aircraft) and concepts of operations for these systems to address the intelligent, increasingly hostile, rapidly changing environments faced by our war fighters. These systems must detect, identify, locate, prioritize, and employ ordinance to achieve permanent destruction of high value targets. New developments in intelligent control, vision based control, mission planning, path planning, decision aiding, communication architectures, logistics and support concepts, and last but not least, software development, validation, and verification are needed to support these systems and make them affordable.

2.3 The Shift to Information-Based Systems

1. Early applications focused on physics; new applications are increasingly information based
 - a. Ubiquitous computing, communications, sensing
 - b. Embedded systems and software
 - c. Examples: router protocols, power control in wireless, supply chains
2. Role of control
 - a. Stability and robustness of interconnections
 - b. Dynamic resource allocation in presence of uncertainty
 - c. Proof by construction programming?
3. Communications will radically change the use of feedback

- a. Increased access to large amounts of information
- b. Distributed, partially asynchronous computation
- c. Many traditional approaches may not work
4. Role of uncertainty is critical (and largely unexplored)
 - a. Control has a unique role to play here
 - b. Vignette introduction

[Note] *Summarize information from aerospace and transportation, information and networks, robotics and intelligent machines.*

Vignette: RoboCup: a testbed for autonomous collaborative behavior in adversarial environments (Raffaello D'Andrea, Cornell University)

RoboCup is an international collection of robotics and artificial intelligence (AI) competitions. The competitions are fully autonomous (no human intervention) head to head games, whose rules are loosely modeled after the human game of soccer; each team must attempt to score more goals than the opponent, subject to well defined rules and regulations (such as size restrictions, collision avoidance, etc.) The three main competitions are known as the Simulation League, the F2000 League, and the F180 League,

The F180 League is played by 6 inch cube robots on a 2 by 3 meter table, and can be augmented by a global vision system; the addition of global vision shifts the emphasis away from object localization and computer vision to collaborative team strategies and aggressive robot maneuvers. In what follows, we will describe Cornell's experience in the F180 League at the 1999 competition in Stockholm, Sweden and the 2000 competition in Melbourne, Australia.

Cornell was the winner of the F180 League in both years that it competed, 1999 and 2000. The team's success can be directly attributed to the adoption of a systems engineering approach to the problem, and by emphasizing system dynamics and control. The systems engineering approach was instrumental in the complete development of a competitive team in only 9 months (for the 1999 competition). Twenty-five students, a mix of first year graduate students and seniors representing computer science, electrical engineering, and mechanical engineering, were able to construct two fully operational teams by effective project management, by being able to capture the system requirements at an early stage, and by being able to cross disciplinary boundaries and communicate among themselves. A hierarchical decomposition was the means by which the problem complexity was rendered tractable; in particular, the system was

decomposed into estimation and prediction, real time trajectory generation and control, and high level strategy.

Estimation and prediction entailed relatively simple concepts from filtering, tools known to most graduate students in the area of control. In particular, smoothing filters for the vision data and feed-forward estimators to cope with system latency were used to provide an accurate and robust assessment of the game state. Trajectory generation and control consisted of a set of primitives that generated feasible robot trajectories; various relaxation techniques were used to generate trajectories that (1) could quickly be computed in real time (typically less than 1000 floating point operations), and (2) took full advantage of the inherent dynamics of the vehicles. In particular, feasible but aggressive trajectories could quickly be generated by solving various relaxations of optimal control problems. These primitives were then used by the high level strategy, essentially a large state-machine.

The high-level strategy was by far the most ad-hoc and heuristic component of the Cornell RoboCup team. The various functions that determined whether passes and interceptions were possible were rigorous, in the sense that they called upon the provably effective trajectory and control primitives, but the high level strategies that determined whether a transition from defense to offense should be made, for example, or what play should be executed, relied heavily on human judgment and observation. As of the writing of this summary, most of the efforts at Cornell have shifted to understanding how the design and verification of high level strategies that respect and fully utilize the system dynamics can take place.

2.4 Opportunities and Challenges Now Facing Us

Although in the past there were often a relevant precedents to serve as guide, the rapid evolution of computation and communication capability has led to a situation in which past solutions are less likely to provide useful models.

Characteristics of the New Environment

The future of Control will be driven by a new environment that differs substantially from that of the past 40 years. Some of the features of this new environment are already apparent and provide insight into the new research directions that must be pursued.

Ubiquitous Computation, Communication and Sensing. The dominant change

in the engineering environment is the presence of ever more powerful computation and cheaper communication. The new software and storage products that these developments have spawned have further changed the engineering landscape in many areas. In addition, microelectronics and MEMS have made available inexpensive sensors whose outputs can be made available via communication networks. It will take decades to take full advantage of these developments. Some innovation will involve standalone, or “receiver” type items, and some will involve extreme interconnectedness of the type seen in the telephone system and its descendants. Both types may, and probably will, depend on the use of feedback control. The new ideas required to be successful in the two cases are, however, likely to be qualitatively different because we do not yet have a great deal of experience in building and operating safe, reliable, highly interconnected systems.

Complexity. Air traffic control systems, power grid control systems and other large-scale, interconnected systems are typical of a class of problems whose complexity is fixed not by the designer but rather by economic considerations and the natural scale of the problem. An acceptable solution in this context must be capable of dealing with the given complexity. In deciding if a system can be built or not, it is important to correctly gauge the feasibility because there is no value in a product that almost works.

Every discipline has methods for dealing with some types of complexity. In the physical sciences, for example, the tools developed for studying statistical mechanics have led to a very substantial body of literature, effective for solving some problems. However, in discussing complexity it is one thing to find a point of view from which aspects of the behavior is compressible (e.g., the entropy of a closed system can only increase”) but it is another to have a “theory of complex systems”. The latter is something of an oxymoron, in that it suggests that the system is not really complex. On the other hand, it does make sense to seek to understand and organize the methodologies which have proven to be useful in the design of highly interconnected systems and to study naturally occurring systems with this in mind. Engineers looking at the immune system may very well be able to suggest new methods to defeat Internet viruses and ideas from neuroscience may inspire new developments in building reliable systems using unreliable components.

Reliable Systems with Unreliable Parts. Most reasonably complex man-made systems are not rendered inoperable by the failure of any particular component and biological systems often demonstrate remarkable robustness in this regard. Simple redundancy, or the spare parts approach to such

problems, is of limited effectiveness because it is uneconomical. Designs that allow the system to reconfigure itself when a component fails, even if this degrades the performance roughly in proportion to the magnitude of the failure, are usually preferred. Although computer memory chips and disk drive controllers often take advantage of strategies of the type, it is still true that the design of self healing systems is not well studied or analyzed.

This issue takes on considerable significance when dealing with interconnected systems of the complexity of the Internet. In this case there are billions of components and yet the system is so essential that little down time can be tolerated.

Disruptive Change. The conventional engineering view of the future is often encapsulated by the words, “better, faster, cheaper”. This evolutionary attitude has resulted in amazing progress and continues to represent an important part of engineering. However, this is not the point of view that will lead to a new air traffic control system or the development of a quantum computer. The recognition of the potential of systems that require and exploit carefully managed complexity has been slow in coming even as the number of truly complex systems integral to our lives is growing rapidly. The systems used for managing the electric power grid or the control of traffic on the Internet are, perhaps, only preliminary examples of what we will see in another decade.

Control has the possibility of producing disruptive change in a number of areas, both technological and scientific. The tools that have been developed over the past 40 years will change the way in which science is pursued and new discoveries are made.

Vignette: Quantum Measurement and Control (Hideo Mabuchi, Caltech)

To illustrate the potentially disruptive power of Control, consider the research of Hideo Mabuchi, who is exploring the use of feedback and control in quantum systems and its implications for unifying quantum and classical physics:

A grand enigma, which is perhaps our primary legacy from 20th Century physics, is that the states and dynamics we ascribe to microscopic (quantum) systems seem incompatible with macroscopic (classical) phenomenology. For example, physical theory claims that it should be illogical simultaneously to assign definite values to certain sets of measurable properties of a quantum system. And yet

we want to believe that quantum mechanics is a correct description of microscopic physics, which evolves robustly into classical dynamics for systems of sufficiently large size and with a sufficiently high degree of interconnection among their manifold degrees of freedom. How can we understand the consistency of quantum mechanics, as a microscopic theory, with classical physics as a manifestly valid description of macroscopic phenomena?

Control theory provides a new set of tools for understanding quantum systems. One set of tools is through systematic techniques for model reduction:

Viewed from a “multiscale” perspective, our challenge in explaining the quantum-classical transition will be to show that classical physics can rigorously be obtained as a robust and parsimonious approximation to the dynamics of certain aggregate degrees of freedom for generic complex quantum systems. In the language of control theory, one would like to derive classical physics as an optimal model reduction of quantum physics. A number of fundamental questions arise as soon as the problem is posed this way. How can this model reduction be so general and robust, depending only upon the structure of quantum theory and not the details of any particular dynamical system? What are the general parameters that control the error bounds on this model reduction? What impact will this program have, if successful, on our basic interpretation of quantum mechanics?

In addition, control can provide new techniques for doing experiments, along us to better explore physical understanding:

... we hope that feedback control will provide a crucial experimental methodology for scrutinizing the validity of quantum measurement theory in realistic laboratory scenarios, especially with regard to the equations for conditional evolution of a system under continuous observation. Such equations could be used as the starting point for controller synthesis, for example, and their validity would be assessed by comparison of experimentally observed closed-loop behavior with theoretical expectations.

Mabuchi’s work illustrates the potential power of control theory as a disruptive technology for understanding the world around us.

Vision for the Future

[Note] *Emphasize the main themes and opportunities of the future. These should be written to directly support the recommendations.*

Approach

[Note] *This section needs to be updated to reflect some of the ideas contained in Chapter 3 and to tie to the recommendations more directly.*

The problems addressed here should be addressed on two fronts. There is a need for a broadly supported, active research program whose goals are to research and further develop methodologies for design and operating reliable and robust highly interactive systems, and there is a need to make room in the academic programs for material specific to this area. The research program should include scientists trained in software engineering, molecular biology, statistical mechanics, systems engineering and psychology. Concrete goals for this work would include the description and rationalization of particular forms of discipline to be used in the design and operation, guidance about the percentage of the resources that should be allocated to security and privacy issues, considerations of the evolvability and expandability of the system, etc. Of course it will be impossible to completely separate these questions from the more obvious “quality of service” issues, so these must be considered at the same time.

One of the characteristics of the Control field has been a high respect for careful thinking, often coupled with an emphasis on clear mathematical formulations of the problems being considered. This discipline has resulted in a body of work that is reliable and unambiguous. Moreover, because this style appeals to some very able graduate students, it has been an important factor in maintaining the flow of talent into the field. However, for engineers and scientists this has been a barrier to entry and can make it difficult for outsiders to assimilate and use the work in their own field. In addition, it has sometimes had a chilling effect on the development of ideas that are not easily translated into mathematics form. The challenge presented by the need to steer a course between the possible extremes here is not new, it has always been present. What is new is that the availability of easily used simulation tools has made the use of heuristic reasoning both more appealing and more reliable. In particular, optimization involving problems that are so large and/or so badly nonconvex that rigorous analysis is infeasible, can now be approached using principled heuristics. Because of the software and computing power now available this may be the most effective way to

proceed. It is important find a place for effective heuristics in the training of students and the highest level professional meetings of the field.

3 Applications, Opportunities, and Challenges

In this chapter, we consider some of the opportunities and challenges for Control in different application areas. These areas are not comprehensive, but represent some of the areas in which Control has been historically important as well as some of the emerging areas that will drive Control theory, technology and practice forward.

[Note] *These sections are still in early draft form. Comments and suggestions are appreciated.*

3.1 Aerospace and Transportation

Aerospace and transportation encompasses a collection of critically important application areas where Control is a key enabling technology. These application areas represent a very large part of the developed world's overall technological capability. They are also a major part of its economic strength, and they contribute greatly to the well being of its people. The historical role of control in these application areas, the current challenges in these areas, and the projected future needs all strongly support the key findings of this report.

The Historical Role

In aerospace, specifically, control has been a key technological capability tracing back to the very beginning of the 20th Century. Indeed, the Wright Brothers are correctly famous not simply for demonstrating powered flight—they actually demonstrated controlled powered flight. Their early Wright Flyer incorporated moving control surfaces (vertical fins and canards) and warpable wings that allowed the pilot to regulate the aircraft's flight. In fact, the aircraft itself was not stable, so continuous pilot corrections were mandatory. This early example of controlled flight is followed by a fascinating success story of continuous improvements in flight control technology, culminating in the very high performance, highly reliable automatic flight control systems we see on modern commercial and military aircraft today [see Flight Control Vignette].

Similar success stories for control technology occurred in many other aerospace application areas. Early World War II bombsights and fire control servo systems have evolved into today's highly accurate radar guided guns and precision guided weapons. Early failure-prone space missions have evolved into routine launch operations, into manned landings on the moon, permanently manned space stations, robotic vehicles roving Mars, orbiting vehicles at the outer planets, and a host of commercial and military satellites serving various surveillance, communication, navigation and earth observation needs.

Similarly, control technology has played a key role in the continuing improvement and evolution of transportation—in our cars, highways, trains, ships and air transportation systems. Control's contribution to the dramatic increases of safety, reliability and fuel economy of the automobile is particularly noteworthy. Cars have advanced from manually tuned mechanical/pneumatic technology to computer controlled operation of all major

functions including fuel injection, emission control, cruise control, braking, cabin comfort, etc. [see Emission Control Vignette]. Indeed, modern automobiles carry dozens (?) of individual processors to see to it that these functions are performed accurately and reliably over long periods of time and in very tough environments.

As a historical note, the cruise control option introduced in the mid-1960s was one of the first servo systems receiving very broad public exposure. Our society's inherent trust in control technology traces back to the success of such early control systems.

Certainly, each of these successes owes its debt to improvements in many technologies, e.g. propulsion, materials, electronics, computers, sensors, navigation instruments, etc. However, they also depend in no small part on the continuous improvements that have occurred over the century in the theory, analysis methods and design tools of Control. As an example, "old timers" in the flight control engineering community still tell the story that early control systems (circa World War II) were designed by manually tuning feedback gains in flight—in essence, trial-and-error design performed on the actual aircraft. Dynamic modeling methods for aircraft were in their infancy at that time, and formal frequency-domain design theories to stabilize and shape single-input single-output feedback loops were still only subjects of academic study. Their incorporation into engineering practice revolutionized the field, enabling successful feedback systems design for ever more complex applications, consistently, with minimal trial-and-error, and with reasonable total engineering effort.

Of course, the formal modeling, analysis and control system design methods described above have advanced dramatically since mid-century. The state of the art today lets us design controllers for much more than single-input single-output systems. The theory and tools handle many inputs, many outputs, complex uncertain dynamic behavior, difficult disturbance environments, and ambitious performance goals. In modern aircraft and transportation vehicles, dozens of feedback loops are not uncommon, and in process control number of loops reaches well into the hundreds. Our ability to design and operate such systems consistently, reliably, and cost effectively rests in large part on the accomplishments of Control over the latter half of the century.

Current Challenges and Future Needs

Still, the control needs of some engineered systems today and those of many in the future outstrip the power of these tools and theories. This is so be-

cause our current tools and theories apply most directly to problems whose dynamic behaviors are smooth and continuous, governed by underlying laws of physics and represented mathematically by (usually large) systems of differential equations. Most of the generality and the rigorously provable features of our methods can be traced to this nature of the underlying dynamics.

Many new control design problems no longer satisfy these underlying characteristics, at least in part. Design problems have grown from so-called “inner loops” in a control hierarchy (e.g. regulating a specified flight parameter) to various “outer loop” functions which provide logical regulation of operating modes, vehicle configurations, payload configurations, health status, etc. For aircraft, these functions are collectively called “vehicle management”. They have historically been performed by pilots or other human operators and have thus fallen on the other side of the man-machine boundary between humans and automation. Today, that boundary is moving!

There are compelling reasons for the boundary to move. They include economics (two, one or no crew members in the cockpit vs. three), safety (no operators exposed to dangerous or hostile environments), and performance (no operator-imposed maneuver limits). A current example of these factors in action is the growing trend in all branches of the military services to field unmanned vehicles. Certain benign uses of such vehicles are already commonplace (e.g. reconnaissance and surveillance), while other more lethal ones are in serious development (e.g. combat UAVs for suppression of enemy air defenses). Control design efforts for such applications must necessarily tackle the entire problem, including the traditional inner loops, the vehicle management functions, and even the higher-level “mission management” functions coordinating groups of vehicles intent on satisfying specified mission objectives.

Today’s engineering methods for designing the upper layers of this hierarchy are far from formal and systematic. In essence, they consist of collecting long lists of logical if-then-else rules from experts, programming these rules, and simulating their execution in operating environments. Because the logical rules provide no inherent smoothness (any state transition is possible) only simulation can be used for evaluation and only exhaustive simulation can guarantee good design properties. Clearly, this is an unacceptable circumstance—one where the strong system-theoretic background and the tradition of rigor held by the Controls community can make substantial contributions.

One can speculate about the forms that improved theories and tools for non-smooth (hybrid) dynamical systems might take. For example, it may

be possible to impose formal restrictions on permitted logical operations, to play a regularizing role comparable to laws of physics. If rigorously obeyed, these restrictions could make resulting systems amenable to formal analyses and proofs of desired properties. This approach is similar to computer language design, and provides support for one of the findings of this report, namely that the Control and Computer Science disciplines need to continue their intimate interactions. It is also likely that our traditional standards of formal rigor must expand to firmly embrace computation, algorithmic solutions, and heuristics.

However, we must not ever lose sight of the key distinguishing features of the Controls discipline, including the need for hard real time execution of control laws and the need for ultra-reliable operation of all hardware and software control components. Many controlled systems today (auto-land systems of commercial transports, launch boosters, F-16 and B-2 aircraft, certain power plants, certain chemical process plants, etc.) fail catastrophically in the event of control hardware failures, and many future systems, including the unmanned vehicles mentioned above, share this property. But the future of aerospace and transportation holds still more complex challenges. We noted above that changes in the underlying dynamics of our control design problems from continuous to hybrid are well under way. An even more dramatic trend on the horizon is a change in dynamics to large collections of distributed entities with local computation, global communication connections, very little regularity imposed by laws of physics, and no possibility to impose centralized control actions. Examples of this trend include the national air space management problem, the automated highway/traffic management problem, and the problem of managing future battlefields.

The national air space problem is particularly significant today, with impending gridlock and congestion threatening the integrity of the existing air transportation system. Many studies are under way attempting to modernize the way traffic is managed, the way individual aircraft schedules and flight paths are established, and the way the system adjusts to upsets due to local weather, local equipment failures, and various other disturbances. General solutions being explored are called “free flight”. They involve distributed calculations of flight plans and trajectories aboard individual aircraft, free of established air corridors, flight plan coordination via negotiations and ground based assistance, and automated collision avoidance technology. This is yet another application where the strong system-theoretic background and the tradition of rigor held by the Controls community can make substantial contributions (see flight management vignette for more details).

Finally, it is important to observe that the future also holds many applications that fall under our traditional control design paradigm, yet are worthy of research support because of their great impact. Conventional “inner loops” in automobiles, but for non-conventional power plants, are examples. Hybrid cars combining electrical drives and low-power internal combustion engines and fuel cell powered cars combining electrical drives with fuel cell generation both depend heavily of well-designed control systems to operate efficiently and reliably. Similarly, increased automation of traditional transportation systems such as ships and railroad cars, with added instrumentation and cargo-tracking systems will rely on advanced controls and schedule optimization to achieve maximum economic impact. Another conventional area is general aviation, where control systems to make small aircraft easy and safe to fly and increased automation to manage them are essential needs.

Other Trends in Aerospace and Transportation

- Aerospace no longer the prime mover?
- Others from workshop ?

3.2 Information and Networks

The rapid growth of communication networks provides several major opportunities and challenges for systems and control. Although there is overlap, we can divide these roughly into two main areas: control of networks, and control over networks.

Control of networks

Control of networks is a large area, spanning many topics, a few of which are briefly described below.

In routing control, the flow of packets through the network is controlled. In the simplest form each router must decide the order, and which output link(s) the packets should be sent to on their way to their final destination(s). Uncertainties include varying link congestion, delays, and rates, and even varying network topology (e.g., a link goes down, or new nodes or links become available), as well as future traffic levels. Current decisions clearly affect the future state of the network, such as the future traffic on links, future buffer levels, queuing delays, etc. Resources that must be managed include router resources like buffer limits, and link resources, such as capacities. Performance is judged in many ways: latency, delay, loss rates, bandwidth, for various streams and other types of traffic. Other measures include how well and how fast the network adapts to changing network congestion, changing traffic patterns, etc.

Several features of these control problems make them very challenging. One is the extremely large scale of the system. Another is the variation in network topology that the routing control must be able handle. Yet another is the decentralized nature of the control: local decisions must be made fast, and so have to be based on locally available information. Of course, information about link congestion, router queues, and traffic demands can be sent over the network between routers. This takes away capacity from real network traffic, and also provides delayed information, which may not be as useful for control. Another complicating issue is the different types and classes of network traffic, all with different requirements for quality of service, in terms of delay, bandwidth, loss probability, etc.

Like many control problems, routing control can be decomposed into several time scales, with very fast decisions made in hardware using lookup tables, which in turn are updated on a slower time scale. At the other extreme in time scale we have optimal network planning, in which new links and nodes are proposed to meet (predicted) rising traffic demand.

Another example of a challenging dynamic resource allocation problem arising in networks is optimal forward caching. The task is to copy documents (or services) that are likely to be accessed often, from many different locations, on multiple servers. When the document is requested, it is returned by the nearest (in network topology sense) server, thereby reducing network traffic and delay. If the source document changes, the changes (at least) must be transmitted to the servers, which consumes some network bandwidth that otherwise would have been available to real network traffic. The problem is to devise a scheme for how often to update, and where to cache copies of documents, based on predictions of access patterns and network congestion.

Several other important control problems arise in wireless networks, including for example the important question of power control. Here the transmitters must decide on an appropriate transmit power, which guarantees a large enough signal to noise and interference ratio (SINR) at the receiver for accurate reception. The subtlety here is that the transmitted signal also appears as an interference term for other wireless links, so increasing the power of a transmitter can affect the SINR of many wireless links. here too, we require fast, possibly asynchronous, somewhat decentralized decisions for a very large, coupled system.

In ad hoc wireless networks, the power control problem is coupled with the routing problem. In the simplest case each of n nodes can transmit to any other node. When nodes are far apart, higher transmit power is required, which in turn has the adverse affect of increasing the interference for other links (thus requiring other transmitters to increase their power). A more sophisticated model connects the transmit powers, or really the receiver SINR, to the maximum possible bit rate (i.e., capacity) over the link. In an ad hoc wireless network, packets can be routed from source to destination over a sequence of nodes. here we seek joint power and routing control algorithms that adapt to rapidly changing network conditions, connectivity, and traffic conditions to satisfy constraints on bandwidth, delay, packet loss, and efficiency (as measured, for example, by total transmit power or total spare capacity). Here too we seek algorithms that are mostly decentralized, or efficiently use internode communication for control coordination.

Control of networks extends beyond data and communication networks. Optimal routing and flow control of commercial aircraft (with emphasis on guaranteeing safe inter-vehicle distances) will help maximize utilization of airports. The (network and software) infrastructure for supply chain systems is being built right now, and simple automated supply chain management systems are beginning to be deployed. In the near future, sophisticated

optimization and control methods can be used to direct the flow of goods and money between suppliers, assemblers and processors, and customers.

Control over networks

As existing networks continue to build out, and network technology becomes cheaper and more reliable than fixed point-to-point connections, even in small localized systems, more and more control systems will operate over networks. We can foresee sensor, actuator, diagnostic, and command and coordination signals all traveling over data networks. The estimation and control functions can be distributed across multiple processors, also linked by data networks. (For example, smart sensors can perform substantial local signal processing before forwarding relevant information over a network.)

Current control systems are almost universally based on synchronous, clocked systems, so they require communications networks that guarantee delivery of sensor, actuator, and other signals with a known, fixed delay. While current control systems are robust to variations that are included in the design process (such as a variation in some aerodynamic coefficient, motor constant, or moment of inertia), they are not at all tolerant of (unmodeled) communication delays, or dropped or lost sensor or actuator packets. Current control system technology is based on a simple communication architecture: all signals travel over synchronous dedicated links, with known (or worst-case bounded) delays, and no packet loss. Small dedicated communication networks can be configured to meet these demanding specifications for control systems, but a very interesting question is:

Can we develop a theory and practice for control systems that operate in a distributed, asynchronous, packet-based environment?

It is very interesting to compare current control system technology with current packet-based data networks. Data networks are extremely robust to gross, unpredicted changes in topology (such as loss of a node or a link); packets are simply re-sent or re-routed to their destination. Data networks are self-configuring: we can add new nodes and links, and soon enough packets are flowing through them. One of the amazing attributes of data networks is that, with good architecture and protocol design, they can be far *more* reliable than their components. This is sharp contrast with modern control systems, which are only as reliable as their weakest link. Robustness to component failure must be designed in, by hand (and is, for safety critical system).

Looking forward, we can imagine a marriage of current control systems and networks. The goal is an architecture, and design and analysis methods, for distributed control systems that operate in a packet-based network. If this is done correctly, we might be able to combine the good qualities of a robust control system, i.e., high performance and robustness to parameter variation and model mismatch, with the good qualities of a network, i.e. self-configuring, robustness to gross topology changes and component failures, and reliability exceeding that of its components.

One can imagine systems where sensors asynchronously burst packets onto the network, control processors process the data and send it out to actuators. Packets can be delayed varying amounts, or even lost. Communication links can go down, or become congested. Sensors and actuators themselves become unavailable or available. New sensors, actuators, and processors can be added to the system, which automatically reconfigures itself to make use of the new resources. As long as there are enough sensors and actuators available, and enough of the packets are getting through, the whole system works (although we imagine not as well as with a dedicated, synchronous control system). This is of course very different from any existing current high performance control system.

It is clear that for some applications, current control methods, based on synchronous clocked systems and networks that guarantee arrival and bound delays for all communications, are the best choice. There is no reason not to configure the controller for a jet engine as it is now, i.e., a synchronous system with guaranteed links between sensors, processors, and actuators. But for consumer applications not requiring the absolute highest performance, the added robustness and self-reconfiguring abilities of a packet-based control system could make up for the lost performance. In any case what will emerge will probably be something in between the two extremes, of a totally synchronous system and a totally asynchronous packet-based system.

Clearly, several fundamental control concepts will not make the transition to an asynchronous, packet-based environment. The most obvious casualty will be the transfer function, and all the other concepts associated with LTI systems (impulse and step response, frequency response, spectrum, bandwidth, etc.). This is not a small loss, as this has been a foundation of control engineering since about 1930. With the loss goes a lot of intuition and understanding. For example, Bode plots were introduced in the 1930s to understand and design feedback amplifiers, were updated to handle discrete-time control systems in the 1960s, and robust MIMO control systems in the 1980s (via singular value plots). Even the optimal control methods in the 1960s, which appeared at first to be quite removed from frequency domain

concepts, were shown to be nicely interpreted via transfer functions.

So what methods will make the transition? Many of the methods related to optimal control, and optimal dynamic resource allocation will likely transpose gracefully to an asynchronous, packet-based environment. A related concept that is likely to survive is also one of the oldest: Lyapunov functions (which were introduced in 1890).

Here is an example of how an “old” concept from control will update gracefully. The idea is that of the Bellman value function, which gives the optimal value of some control problem, posed as an optimization problem, as a function of the starting state. It was studied by Pontryagin and other pioneers of optimal control in the 1940s, and has recently had a resurgence (in generalized form) under the name of control-Lyapunov function. It is a key concept in dynamic programming.

The basic idea of a control-Lyapunov function (or the Bellman value function) is this: If you knew the function, then the best thing to do is to choose current actions that minimize the value function in the current step, without any regard for future effects. (In other words, we ignore the dynamics of the system.) By doing this we are actually carrying out an optimal control for the problem. In other words, the value function is the cost function whose greedy minimization actually yields the optimal control for the original problem, taking the system dynamics into account. In the work of the 1950s and 60s, the value function is just a mathematical stepping stone toward the solution of optimal control problems.

But the idea of value function transposes to an asynchronous system very nicely. If the value function, or some approximation, were broadcast to the actuators, then each actuator could take independent and separate action, i.e., each would do whatever it could to decrease the value function. If the actuator were unavailable, then it would do nothing. In general the actions of multiple actuators has to be carefully coordinated; simple examples show that turning on two feedback systems, each with its own sensor and actuator, simultaneously, can lead to disastrous loss of performance, or even instability. But if there is a value or control Lyapunov function that each is separately minimizing, everything is fine; the actions are automatically coordinated (via the value function).

Another idea that will gracefully extend to asynchronous packet-based control is model predictive control. The basic idea is to carry out far more computation at run time, by solving optimization problems in the real-time feedback control law. Model predictive control has played a major role in process control, and also in supply-chain management, but not (yet) in other areas, mainly owing to the very large computational burden it places

on the controller implementation. The idea is very simple: at each time step we formulate the optimal control problem, up to some time horizon in the future, and solve for the whole optimal trajectory (say, using quadratic programming). We then use the current optimal input as the actuator signal. We use the sensor signals to update the model, and carry the same process out again.

3.3 Robotics and Intelligent Machines

Overview

The goal of cybernetic engineering, already articulated in the 1940s and even before, has been that of implementing systems capable of exhibiting highly flexible or “intelligent” responses to changing circumstances. In [1] the MIT mathematician Norbert Wiener, gave a widely read, albeit completely non-mathematical, account. The basic electromechanical engineering and the computing capabilities required to build practical systems of this type have evolved over the last half-century to the point where today there exists rapidly expanding possibilities for making progress toward these long held goals. The implementation of principled and moderately sophisticated algorithms is already possible on available computing hardware and more capability will be here soon. The successful demonstrations of vision guided automobiles operating at high speed, the use of robotic devices in manufacturing, and the commercialization of mobile robotic devices, attest to the practicality of this field. Important problems under investigation include questions about the use of high data rate sensors, large data bases, real-time distributed control and theories of adaptation and learning.

Although early demonstrations involving drone airplanes and pilotless vehicles were often little more than remotely controlled servomechanisms and the most visible activity in the field of intelligent machines now centers around robotic hardware, it requires little imagination to see that these examples will eventually give way to more sophisticated types of intelligent machines. We already see the signs in new devices such as the Sony dog and its peers. It seems likely that some considerably expanded view of intelligent machine will make use of sensors and feedback control in a more integral way, going well beyond the idea of a robot as a simple position controlled device. A important new idea here, one that is central to the development of the modules used in intelligent machines, is that of a symbol-to-signal transducer. By this we mean components which accept symbol strings as inputs and produce analog signals or motions as outputs. This point of view places the field in firm contact with computer science and communications engineering.

Robotic Engineering: Control and Software

Robotics is a broad field; the perspectives afforded by computer science, electrical engineering, mechanical engineering, psychology, and neuroscience all yield important insights. Even so, there are pervasive common threads,

such as the understanding and control of spatial relations and their time evolution. The emergence of the field of robotics has provided the occasion to analyze, and to attempt to replicate, the patterns of movement required to accomplish useful tasks. On the whole, this has been sobering experience. Just as the ever closer examination of the physical world occasionally reveals inadequacies in our vocabulary and mathematics, roboticists have found that it is quite awkward to give precise, succinct descriptions of effective movements using the syntax and semantics in common use. Because the motion generated by a robot is usually its *raison d'être*, it is logical to regard motion control as being the central problem. Its study has raised several new questions for the control engineer relating to the major themes of feedback, stability, optimization and estimation. For example, at what level of detail in modeling (i.e. kinematic or dynamic, linear or nonlinear, deterministic or stochastic, etc.) does optimization enter in a meaningful way? Questions of coordination, sensitivity reduction, stability, etc. all arise.

In addition to these themes, there is the development of software. At present there is almost no transportability of robotic motion control languages. The idea of vendor independent languages that apply with no change to a wide range of computing platforms and peripherals has not yet been made to work in the field of robotics. The clear success of such notions when applied to operating systems, languages, networks, disk drives and printers makes it clear that this is a major stumbling block. What is missing is a consensus about how one should structure and standardize a “motion description language”. Such a language should, in addition to other things, allow one to implement compliance control in a general and natural way.

From the point of view of the programmer, high level languages are more efficient than low level ones. Likewise in directing the motion of a robot the programmer would like to say as little as possible about the means, preferring to focus attention on the ends. In order to make this possible, it is necessary to incorporate automatic path planning algorithms in the software and to write compilers that are capable of converting high level directives into the motor control programs needed to execute motion segments. This means that motion planning algorithms are an important part of any high level programming environment and that proving correctness of the motion programs produced by such compilers is an issue.

The contrast between the sophistication and ease of use of computer programs written for everyday applications such as word processing, graphics editing and CAD/CAM, little is available to help generate the description of a motion which is to be executed by a robot. Producing such software will not be an easy task. Just as the step from text processing to graphics

is made difficult by the two dimensionality of graphics, as opposed to the one dimensionality of a string of text, motion generation will be intrinsically more difficult than graphics by virtue of the fact that it is three dimensional and it evolves in time. Clues that can be gleaned from the study of motion generation for animation may provide guidance as one attempts to put in place the data structures, algorithms and database systems required to support an improved programming environment. In another vein, the use of computers for real time control is currently driving computer technology as hard or harder than any other application area but it has seemed difficult to formulate a general, theoretically sound, approach to this class of problems. It has proven to be difficult to use the inspiration provided by biological systems to generate designs suitable for implementation with today's computer hardware.

A Nexus of Communication, Computation and Control

Although for several decades it has been felt that computation, communication and control are the basic methodologies for the nonphysical aspects of electrical engineering, there have been rather few examples of applications in which the use of these disciplines is truly integrated. Up until now they have coexisted in many systems but they usually do not interact in any very significant way. This is a result of technological limitations. The most demanding applications of automatic control have tended to involve feedback control of dynamical systems. In such cases time is critical. Until recently the computational hardware required to achieve real-time control have been expensive and difficult to use. Only a few specialized applications have received the attention required to implement a close coupling between computing and control. Communication theory has been very successful in developing appropriate strategies for dealing with uncertain channels, in applying digital methods to improve fidelity and in using coding theory to store and retrieve data with great speed and reliability. Those areas which appear to have the most potential for use in control problems, such as algorithmic approaches to rate distortion trade offs and coding to a fidelity criterion, have only come in to their own rather recently. (See [4].)

Certain robotic applications, such as those that call for the use of vision systems to guide robots, now require the use of computing, communication and control in an integrated way. The computing that is to be done must be opportunistic, i.e. it must be tailored to fit the needs of the specific situation being encountered. The data compression that is needed to transmit television signals to a computer must be done with a view toward how the

results will be used by the control system. It is both technologically difficult and potentially dangerous to build complex systems that are controlled in a completely centralized way. For this reason we need to decide how to distribute the control function over the communication system. Recent work on the theory of communication protocols has made available better methods for designing efficient distributed algorithms; perhaps this work can be adapted in such a way as to serve the needs of robotic applications.

Recently there have begun to appear commercial devices that provide reasonably wide-band wireless local area communication; they have enormous potential for simplifying and enhancing the operation of mobile robots. These systems can be divided into those that are free-space, line-of-sight systems versus free-space omni-directional systems. The advantages of the omni-directional systems based on radio links are that they do not require pointing control. The advantage of the directed beam systems based on infrared transmission is that they are relative immune to cross-talk and interference and they do not flood the area with their signal. Because of this technology one can now build tetherless vision guided robots with remote computational support. In effect, wireless communication allows the experimenter to give brains to mobile robots and legs to computational engines.

Control using High Data-Rate Sensors

Without large expenditure, we are able to gather and store more pictures and sounds, temperatures and particle counts, than we know how to use. Yet we continue to witness occasional catastrophic failures of our man-machine systems, such as those used for transportation, because we do not correctly interpret, or appropriately act on, the information available to us. It is apparent that in many situations, collecting the information is the easy part. Feedback control embodies the idea that performance can be improved by coupling measurement directly to action. Physiology provides many examples attesting to the effectiveness of this technique. However, as engineers and scientists turn their attention to the highly automated systems currently being built by the more advanced manufacturing and service industries, they often find that the direct application of feedback control is frustrated by a web of interactions which make the smallest conceptual unit too complex for the usual type of analysis. In particular, vision guided systems are difficult to design and often fail to be robust with respect to lighting conditions and changes in the environment. In order to proceed, it seems, design and performance evaluation must make more explicit use of ideas such as adaptation, self-diagnosis and self-optimization.

Indications are that the solution to the problems raised above will involve active feedback control of the perceptual processes. One area that has received considerable attention is the area of active vision in which the vision sensor is controlled on the basis of the data it generates. Other work involves tuning the vision processing algorithms on basis of the data collected. The significant progress now being made toward the resolution of some of the basic problems results, in large part, from the discovery and aggressive use of highly nonlinear signal processing techniques. Examples include the variational theories that have been brought to bear on the image segmentation problem, the theories of learning based on computational complexity, and information theoretic based approaches to perceptual problems. Attempts to incorporate perceptual modules into larger systems, however, often raise problems about communication and distributed computation which are not yet solved. (See [3].)

Linkages

Through the study of robotics the field of control reaches out and links with a number of other important disciplines. We will discuss three: Software, Mechanics and Psychology.

Software. The use of computers for real time control is currently driving computer technology as hard or harder than any other application area but it has seemed difficult to formulate a general, theoretically sound, approach to this class of problems. Perhaps these questions are best studied with one hand on a neuroscience book and the other on a bus architecture manual.

Many of the important algorithmic issues which must be confronted in the implementation of robotic systems involve the change of coordinates necessary to permit an efficient description of relative position. Important questions include: How should a computer carry out coordinate changes? What methods are best for computing inverse kinematics? Is it efficient to implement specific differentiable manifolds as data types? What are economical but adequate sets of coordinate systems for motion control?

In the area of software development there are also important questions. One goal is to understand spatial description languages at a sufficiently deep level so as to be able to give an efficient descriptions of a part which are being machined by a milling machine which are, at the same time, adapted to work well with the algorithms which keep the machine tool from colliding with the part being machined or to be able to describe the free space in a cluttered room efficiently enough so that a mobile robot can move around with running in to things.

It is also worth pointing out that the field of robotics, being the most active focal point for the interaction between computer science and control, has become an important testbed for the development of the whole field of automatic control. It seems clear that more extensive use of computer control, be it for factories, automobiles or homes, will be most effective if it comes with a natural human interface. Having this goal in mind, one should look for interfaces which are not only suitable for the given application but which are sufficiently general so that with minor modification, they can serve in related applications as well. One may argue that biologically inspired solutions deserve serious consideration in view of their proven robustness and manifest naturality.

Mechanics. The concept of a kinematic chain is basic to robotic manipulation and these objects show up in considerable variety in practical applications. It is, therefore, particularly pleasant to observe that a very natural description of kinematic chains is afforded by the mathematical ideas found in nonlinear control. It turns out that a key step in the design of controllers for industrial robots is equivalent to finding an algorithm for converting between coordinates of the first and second type for the group of rigid motions in three dimensions. We mention a second, perhaps unexpected, mathematical fact related to the manipulation of objects. In considering the application of grasping forces to objects, systems of inequalities play a central role because fingers can only push against, and not pull, objects. In fact, the study of grasping involves convex analysis, models of friction and details of the interface between the hand and the object that go considerably beyond simple mechanics. It happens that many tasks, including, but not limited to, grasping, can be done in more than one way. This may happen because the robot has more than the minimal number of degrees of freedom or because the task description is ambiguous. Disposing of such problems is usually called resolution of redundancy and in some cases it leads to non-local questions in geometry.

Psychology. Body movement is the most obvious and easily measured manifestation of brain activity, (“The mind exists to control the body.”) However, the effects of the brain on motion are confounded with effects from almost all other systems. This has discouraged work on the role of the brain in motion control, contributing heavily to a state of affairs in which other subjects, such as vision, have pushed far ahead in terms of the detailed understanding of neural processing. Interesting questions include:

1. What is the nature of the motion control languages found in animals? For example, what kind of extensibility is to be found in nature’s

motion control languages?

2. How many changes of coordinates are there and where are they computed in the mammalian brain?
3. What is the mathematical expression of mental practice? How can the brain deliver a better control sequence as a result of reviewing the motion in abstract form?

It is generally thought that the premotor cortex generates a plan for motor behavior and that this plan is transmitted through the motor cortex for execution. What resources are used by the premotor cortex in this planning process? It is to be expected that the synthesis occurring in the premotor cortex would draw on a data base of previously used motion segments together with a motion sequence editor which would provide the various types of geometric transformations, scaling and time warping necessary to adapt these motion segments to the purpose at hand. It is also likely that there is some type of previewing of the motion, possibly involving visual imagery and feedback if the previewed image suggests any difficulties. We are about to begin to write programs intended to simplify the process of choreographing robotic movement. When they reach maturity, such "premotor cortex" software can be expected to be the word processor/page-layout programs of the robotics world and hence quite central to its development.

Conclusion and Epilogue

In the longer run, it is expected that there will evolve an effective set of concepts and principles which will allow us to organize and simplify what we now know as isolated facts about intelligent machines. As to the possibility that these general principles will be best expressed in terms of ideas now available or perhaps in some other form is, of course, not known. None the less, the goal of this research endeavor is to develop and articulate such principles. The field of robotics is, perhaps, the most active focal point for the interaction between computer science and control, and has become an important test bed for the development of the whole field of automatic control. It seems clear that more extensive use of computer control, be it for factories, automobiles or homes, will be most effective if it comes with a natural human interface. Having this goal in mind, one should look for interfaces which are not only suitable for the given application but which are sufficiently general so that, with minor modification, they can serve in related applications as well. One may argue that biologically inspired

solutions deserve serious consideration in view of their proven robustness and manifest naturalness. It may seem ironic that even for the most advanced robots, motion control has not been developed to the same degree as some highly specialized topics such as the control of printers. However, this is consistent with the relative popularity of soft and hard automation, and reflects the current unavailability of effective, sensory rich and mechanically elegant, cost effective robots.

We have suggested a number of questions here whose answers could hasten the development of new applications of robotics. It seems that the future of robotics is bright, but that considerable patience and hard work is required. Clearly the further development and effective use of smaller and less obtrusive sensors and actuators will be necessary. Just as we have come to base our research agenda on the hypothesis that there will be further developments in microelectronics, we should assume that further developments in micromechanics will lead to practical implementation of ideas that are theoretically sound.

Graduate students will want to know where their work fits in and which departments in the outside world will want to hire them. Personally I am comfortable with the idea that the answer could be different for different people, electrical engineering, computer science, mechanical engineering, applied mathematics and psychology being possible targets.

[Note] *The remaining sections are in the processing of being written and will be included when available.*

3.4 Biology and Medicine

Life sciences?

3.5 Materials and Processing

3.6 Other Applications

Environment

Economics and Finance

4 Education and Outreach

IV. Education

- A. Applications of control are broader than ever
- B. Need to make controls more accessible
- C. Need to provide a broader education for our students
 - 1. Broader grasp of engineering, science, and math disciplines
 - 2. Increased leadership and communications skills

5 Recommendations

[Note] *This final chapter should pull together the ideas of the rest of the report and present some actionable recommendations, broken into the six areas listed below. Pull out main recommendations as boxed comments, with supporting text around them (not yet written).*

5.1 Integrated Control, Computation, Communications, and Networking

Substantially increase research aimed at the *integration* of control, computer science, communications, and networking.

5.2 Control at Higher Levels of Abstraction

Substantially increase research in Control at higher levels of abstraction, moving toward enterprise level systems.

5.3 High-Risk, Long-Range Applications of Control

Explore high-risk, long-range applications of Control to areas such as nanotechnology, quantum mechanics, biology, and environmental science

5.4 Support for Theory and Interaction with Mathematics

Maintain support for theory and interaction with mathematics, broadly interpreted (including areas such as dynamical systems, graph theory, combinatorics, complexity theory, queuing theory, etc).

5.5 New Approaches to Education and Outreach

Invest in new approaches to education and outreach for the dissemination of basic ideas to non-traditional audiences.

5.6 Additional Considerations